

Photometric Stereo for 3D face reconstruction using non linear illumination models

B. Villarini, A. Gkelias, and V. Argyriou

University of Westminster, Imperial College, Kingston University
B.Villarini@westminster.ac.uk, a.gkelias@imperial.ac.uk,
Vasileios.Argyriou@kingston.ac.uk

Abstract. Face recognition in presence of illumination changes, variant pose and different facial expressions is a challenging problem. In this paper, a method for 3D face reconstruction using photometric stereo and without knowing the illumination directions and facial expression is proposed in order to achieve improvement in face recognition. A dimensionality reduction method was introduced to represent the face deformations due to illumination variations and self shadows in a lower space. The obtained mapping function was used to determine the illumination direction of each input image and that direction was used to apply photometric stereo. Experiments with faces were performed in order to evaluate the performance of the proposed scheme. From the experiments it was shown that the proposed approach results very accurate 3D surfaces without knowing the light directions and with a very small differences compared to the case of known directions. As a result the proposed approach is more general and requires less restrictions enabling 3D face recognition methods to operate with less data.

Keywords: face reconstruction, face recognition, photometric stereo, 3D imaging, non-linear dimensionality reduction, illumination models

1 Introduction

Automatic face and facial expression recognition has become a very vibrant topic in the last decade due to the fast progress of human computer intelligent interaction (HCII). Face recognition on frontal faces under controlled condition, such as known light condition, is a mature research field and high recognition accuracy can be achieved. However, the performance decreases in presence of illuminations changes, pose variations, facial expressions and a large number of subjects [12]. A large amount of interest has been addressed towards 3D modeling and reconstruction of faces in order to improve face and facial expression recognition. Photometric stereo techniques are used to estimate the illumination condition and to extract 3D geometry information of a face [17, 18, 2].

Photometric stereo techniques are further categorised into constrained and unconstrained. Unconstrained photometric stereo means that we do not have any priori knowledge of the light-source direction or the light-source intensity. However, in constrained photometric stereo the light-source direction is not an

unknown, and for this reason it is easier to determine the surface normal and the surface reflectance. The main issue related to the unconstrained photometric stereo approaches is that a large number of images is required captured while the light source is moving around the observed object. Therefore, this process is not convenient for real time applications since both the capturing and processing time increases significantly. In this paper in order to solve these problems an unconstrained photometric stereo approach for face reconstruction is proposed requiring at least three images of the surface captured under different unknown illumination directions. The primary contributions of this work are twofold. First, we present a novel methodology for illumination direction estimation for human faces based on low dimensional subspaces. Secondly, we perform an analysis for the accuracy of the obtained lighting conditions in terms of reconstruction accuracy and computational complexity.

This paper is organised as follows. In Section 2 previous work on photometric stereo and low dimensional subspaces is reviewed. In Section 3, we present an overview of the proposed approach. In Section 4 experiments are performed using different datasets and metrics indicating the advantages of the proposed approach. Finally, conclusions on the proposed methodology and the evaluation process are presented.

2 Previous work

Woodham [30] was the first to introduce photometric stereo. He proposed a method which was simple and efficient, but only dealt with Lambertian surfaces and was sensitive to noise. An unconstrained photometric-stereo method for estimating the surface normal and the surface reflectance of objects without a priori knowledge of the light source directions was proposed in [7]. Also, worth mentioning the work in [16, 9, 29] that are based on similar approaches. Recently approaches for 3D reconstruction based on photometric stereo with unknown lighting were proposed [4, 25, 1, 26, 31, 22, 21, 27]. Regarding all these methods, the main difference is that they require a significant amount of images captured while the illumination direction is changing, which makes them unsuitable for real time applications and scenarios that involve humans due to their unconstrained self-movement.

To the best of our knowledge, few approaches based on low dimensional subspaces have been used on photometrics. However, a few of them have potential since they are able to consider explicitly the modelling of illumination changes in their methodology. Georghiades et al. [9] a generative model is created by using low dimensional linear subspaces in order to reconstruct new poses and illuminations. Although useful for face recognition, the systems is face specific and therefore unable to extrapolate to new subjects as required in photometrics. Hallinan [10] propose a low-dimensional model for human faces that can both synthesize a face image when given lighting conditions and can estimate lighting conditions when given a face image. Although based on PCA, the method is designed to explain lighting conditions and not to discount them. However this is achieved by limiting all the other sources of variability, including the usage of different people, and will fail otherwise. Lee [20] also propose to represent the

illumination cone in a low dimensional space, similarly to [9], by using spherical harmonics. Again, the generated low dimensional spaces are person specific.

As a common limitation factor, all these approaches rely on linear methodologies, and therefore, on the assumption of linearity in the face subspace. This assumption is only true under certain conditions, such as face specific applications, as we will explain in section 3.1, which lead to poor performances for realistic scenarios with multiple subjects, different illumination conditions and facial expressions. Instead, we propose a non linear low dimensional space able to model these factors and take advantage of it for 3D face reconstruction and related recognition applications.

3 Proposed methodology

In this work a novel two step approach for face reconstruction using photometric stereo is proposed. Initially, illumination direction is estimated using low dimensional illumination models. During this step a large number of faces illuminated from all possible directions on a hemisphere is used for training and to create a manifold that represents all the lighting directions in a 3D space. In the second part of the proposed methodology, a new face is captured and at least 3 images are obtained illuminated from different unknown directions. The images are transferred to the new space using the mapping function obtained from the previous step and the corresponding illumination direction is obtained. Finally, photometric stereo is applied and the surface normals are estimated. The proposed algorithm is analysed in the following sections.

3.1 Illumination subspace

Using a set of training images from different people and lighting directions, a low dimensional subspace can be generated [9]. In this subspace, the distribution of the training samples allows us to estimate the illumination direction while the usage of different subjects allow generalisation to new test subjects.

Dimensionality reduction techniques have been used frequently to model face appearance and pose [11], but illumination variations have been rarely considered or accepted as a factor to be removed [9, 10, 20]. This is mainly due to the false assumption that the low dimensional subspace comprising the face samples is linear. Belhumeur et al. [5] demonstrated that all images of a given Lambertian surfaces, taken from a fixed point of view, and under varying illuminations can lie in a 3D linear subspace, but they also pointed that shadowing, facial expressions and other factors produce that regions of the face may exhibit deviation from a linear subspace. The required usage of several subjects in the dataset, and therefore different Lambertian surfaces, implies a substantial factor that affects this linearity.

In order to deal with all the possible, we propose a nonlinear dimensionality reduction (DR) technique, as oppose to other approaches which try to linearise the space by choosing an ad-hoc optimal linear and discarding non linear factors. Many different non linear DR techniques could be used, both mapping based (GPLVM [19], GPDM [14]) or spectral based (LE [6], LLE [24], Isomap [13]).

However, the big differences among the face surfaces can hide the illumination information, being discarded during the DR process by conventional techniques. Among these methodologies, t-Stochastic Neighbour Embedding (t-SNE) [28] has been proposed to provide a mathematical framework where new constraints can easily be introduced.

3.1.1 Illumination manifold using Stochastic Neighbour Embedding

t-Stochastic Neighbour Embedding (t-SNE) [28] is a non-linear dimensionality reduction technique used to embed high-dimensional data into a low-dimensional space (e.g., two or three dimensions for human-intuitive visualization). Given a set of N high-dimensional faces of people under different illumination conditions (i.e. data-points) x_1, \dots, x_N , t-SNE starts by converting the high-dimensional Euclidean distances between data-points ($\|x_i - x_j\|$) into pairwise similarities given by symmetrized conditional probabilities. In particular, the similarity between data-points x_i and x_j is calculated from (1) as:

$$p_{ij} = \frac{p_{i|j} + p_{j|i}}{2N} \quad (1)$$

where $p_{i|j}$ is the conditional probability that x_i will choose x_j as its neighbour if neighbours were picked in proportion to their probability density under a Gaussian centred at x_i with variance σ_i^2 , given by (2):

$$p_{i|j} = \frac{\exp(-\|x_i - x_j\|^2/2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|x_k - x_i\|^2/2\sigma_i^2)} \quad (2)$$

In the low-dimensional space the Student-t distribution (with a single degree of freedom: $f(x) = 1/(\pi(1+x^2))$) that has much heavier tails than a Gaussian (in order to allow dissimilar objects to be modelled far apart in the map) is used to convert distances into joint probabilities. Therefore, the joint probabilities q_{ij} for the low-dimensional counterparts y_i and y_j of the high-dimensional points x_i and x_j are given by

$$q_{ij} = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}}. \quad (3)$$

The objective of the embedding is to match these two distributions (i.e., (1) and (2)), as well as possible. This can be achieved by minimizing a cost function which is the Kullback-Leibler divergence between the original (p_{ij}) and the induced (q_{ij}) distributions over neighbours for each object

$$D_{KL}(P||Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (4)$$

The minimization of the cost function is performed using a gradient decent method which have the following simple form:

$$\frac{\delta D_{KL}}{\delta y_i} = 4 \sum_j \frac{(p_{ij} - q_{ij})(y_i - y_j)}{(1 + \|y_i - y_j\|^2)} \quad (5)$$

3.1.2 Mapping functions Once the low dimensional illumination space is obtained, a mapping mechanism is needed in order to project a new face or a subset of them and estimate the most likely illumination direction. Methods such as t-SNE allow unsupervised generation of embedded spaces, but they do not provide explicitly any mapping mechanism between the low and high dimensional spaces. This issue has been tackled very effectively by Radial Basis Function Networks (RBFN) [8]. Projection functions are produced by training direct ϕ and inverse ϕ' sets of functions between high and low dimensional spaces.

$$\phi : \mathbb{R}^N \rightarrow \mathbb{R}^n \text{ and } \phi' : \mathbb{R}^n \rightarrow \mathbb{R}^N \quad (6)$$

In our framework, multi-dimensional Gaussian activation functions ϕ_j (Eq. 7) are employed because of its flexibility and superior performance to fit the subspace.

$$\phi_j = e^{-(X-\mu_j)^T \cdot \Sigma_j^{-1} \cdot (X-\mu_j)} \quad (7)$$

for $j = 1, \dots, n_g$, where X is the input feature vector, n_g the number of Gaussian functions to be discovered and μ_j and Σ_j the mean and covariance respectively of each Gaussian function.

3.1.3 Illumination estimation In order to estimate the most likely illumination directions for a new face sample, a nearest neighbour classifier is used in the embedded space. The embedded space is composed of 9 different subspaces, each of them comprising all the possible azimuth angles from 0 to 360, at 9 different elevation angles from 10 to 90 degrees. Each subspace generated with t-SNE produces a radial manifold, where people are overlapping and distributed according to both their appearance and the illumination angle of the lighting. The overall embedded space can be represented as a hemisphere composed of radial manifolds (see figure 1).

By projecting a new sample in this space and pairing with the nearest neighbour, the illumination direction is estimated as the same belonging to this nearest neighbour. This strategy has been proved enough for providing a sufficient estimation, as depicted in figure 1 and in the result section.

3.2 Photometric Stereo without illumination information

Since a mapping function from the high dimensional space of all possible shadow and highlight deformations that could occur on a human face to a 3D space is obtained, the illumination direction for any face could be estimated using the same mapping mechanism. Regarding the illumination deformations over a face and in general a surface, it is well known that the fraction of light reflected on an object's surface in a certain direction depends upon the optical properties of the surface material. In this paper we use the Lambertian model, thus the fraction of the incident illumination reflected in a particular direction depends only on the surface normals.

All faces share common surface characteristics, which therefore result in similar statistical distributions of normal vectors, and therefore the shadow and




Fig. 1. Left. The obtained 3-dimensional manifold using t-SNE corresponding to the hemispherical embedded space for azimuth and elevation angles. Right. Angular error obtained for this sequence using nearest neighbour classifier for subject 3 projected on illumination subspace composed by subjects [1,2, 4-7].

highlight deformations share common characteristics too, independent of the human face. This approach can be extended to any class of surface, not only faces, as long as they share similar facet normal distributions.

Based on the Lambertian model that is used, if \mathbf{n} is the normal vector of a surface facet, ρ its albedo with the cosine of the incidence angle θ_i (the angle between the direction of the incident light and the surface normal), \mathbf{L} is the light direction and I the corresponding brightness value recorded for that facet, we have

$$I = \rho \cos(\theta_i) = \rho(\mathbf{L} \cdot \mathbf{n}) \quad (8)$$

Let us now consider a Lambertian surface patch with albedo ρ and normal \mathbf{n} , illuminated in turn by several fixed and known illumination sources with directions $\mathbf{L}^1, \mathbf{L}^2, \dots, \mathbf{L}^{\tilde{N}}$. In this case we can express the intensities of the obtained pixels as:

$$I^k = \rho(\mathbf{L}^k \cdot \mathbf{n}), \quad \text{where } k = 1, 2, \dots, \tilde{N}. \quad (9)$$

If we move to a matrix form, equation (9) could then be rewritten as

$$\mathbf{I} = \rho[\mathbf{L}]\mathbf{n} \quad (10)$$

If there are at least three illumination vectors which are not coplanar, we can calculate ρ and \mathbf{n} using the Least Squares Error technique, which amounts to applying the left inverse of $[\mathbf{L}]$:

$$([\mathbf{L}]^T[\mathbf{L}])^{-1}[\mathbf{L}]^T\mathbf{I} = \rho\mathbf{n} \quad (11)$$

Since \mathbf{n} has unit length, we can estimate both the surface normal (as the direction of the obtained vector) and the albedo (as its length). Extra images allow one to recover the surface parameters more robustly.

The problem of reconstructing the surface from the modified normals is considered next and the depth map needs to be obtained. Therefore, the surface is

represented as $(x, y, f(x, y))$, and the normal as a function of (x, y) is

$$\mathbf{N}(x, y) = \frac{1}{\sqrt{1 + \frac{\partial f^2}{\partial x} + \frac{\partial f^2}{\partial y}}} \left(-\frac{\partial f}{\partial x}, -\frac{\partial f}{\partial y}, 1 \right)^T \quad (12)$$

To recover the depth map, we need to determine $f(x, y)$ from measured values of the unit normal. There are a number of ways in which a surface may be recovered from a field of surface normals [23, 15]. There are local and global methods based on trigonometry and the minimisation of error functionals, respectively and the most suitable could be selected for this part of the process.

Assume that the measured value of the unit normal at some point (x, y) is $(a(x, y), b(x, y), c(x, y))$. Then

$$\frac{\partial f}{\partial x} = \frac{a(x, y)}{c(x, y)} \quad \frac{\partial f}{\partial y} = \frac{b(x, y)}{c(x, y)} \quad (13)$$

At this stage we may perform another check on our data set. Because

$$\frac{\partial^2 f}{\partial x \partial y} = \frac{\partial^2 f}{\partial y \partial x} \quad (14)$$

we expect

$$A(x, y) \equiv \frac{\partial \left(\frac{a(x, y)}{c(x, y)} \right)}{\partial y} - \frac{\partial \left(\frac{b(x, y)}{c(x, y)} \right)}{\partial x} \quad (15)$$

to be small (close to zero) at each point.

Assuming that the partial derivatives pass the above sanity test, we can reconstruct the surface up to some constant error in depth. The partial derivatives give the change in surface height with a small step in either the x or the y direction. This means that we can get the surface by summing these changes in height along some path. In particular, we have

$$f(x, y) = \oint_C \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right) \cdot d\mathbf{l} + c \quad (16)$$

where C is a curve starting at some fixed point and ending at (x, y) , $d\mathbf{l}$ is the infinitesimal element along the curve and c is a constant of integration, which represents the unknown height of the surface at the starting point. In order to improve further the reconstruction we combine it with a multigrid 2D integration algorithm, which iteratively solves a global minimization problem, and is less sensitive to the propagation of local errors.

4 Experiments and Results

In order to evaluate the accuracy of the proposed methodology a set of experiments was performed using two different datasets and metrics. In more details the first database (see figure 2) was captured in our lab and contains 7 faces male

and female all of them facing the camera. The second dataset (see figure 2) that was used is the Photometric Database presented in [3] captured under a similar setup having more than 300 faces. In both cases, the persons are assumed to be still during the acquisition stage since a high speed camera was used for the acquisition (i.e. 200 frames per second), eliminating the registration problem. In the second stage of the evaluation procedure, photometric stereo is applied on the input images of the databases and the 3D surface of each face is obtained. The faces are aligned using manually placed markers and an affine transformation algorithm. Since the 3D faces are aligned a Lambertian model is used to generate 2D images and shadow maps illuminated under all possible directions (all the possible azimuth angles from 0 to 360, every 10 degrees, and elevation angles from 10 to 90 degrees, every 10 degrees) on a hemisphere as it is shown in figure 3. These images at a resolution of 144×144 are using as input at the t-SNE as a training set to generate a embedded low dimensional space representation of the illumination variations and shadow deformations on a human face in a low dimensional space. Consecutively, using the one leave out process, where each subject is removed from the training set dataset and used for testing, the faces are reconstructed using the estimated illumination directions for the input images.

4.1 Illumination estimation

By using the previously mentioned leave-one-out schema, the performance of the automatic illumination estimation can be evaluated as well as the capability of the embedded space to generalise to new subjects out of the training dataset. Results of the average angular error for each of the datasets are reported in table 1. Regarding the average angular error it was expected to be in that range since two consecutive illumination sources are in range of 15 degrees, which explains the results indicating that the estimated direction is always in the second-order neighborhood. A particular case for subject 3 is depicted in figure 1.

Table 1. Overall performance of the angular estimation provided by the nearest neighbour classifier on the illumination embedded space

	Dataset 1	Dataset 2 [3]
Angular error [Degrees]	32.7536	33.5571

It can be concluded that the illumination embedded space is able to provide a reasonably accurate estimation of the lighting direction, which can be now feed into the reconstruction module.

4.2 Reconstruction performance

In order to evaluate the performance the average difference of the real and the estimated heightmaps was used and furthermore the Hausdorff distance was used to compare the reconstructions with the original profiles.



Fig. 2. (Left) An example of faces part of the first dataset. (Right) An example of faces part of the second dataset [3].




Fig. 3. Examples of possible illumination directions. Each dot on the hemisphere corresponds to a possible light l . It may be identified by its azimuth angle φ_l and its zenith angle θ_l .




Fig. 4. Reconstructed surfaces using the proposed method with unknown illumination on the right and the ground truth on the left under different view points for the first dataset.



Fig. 5. The error difference between the reconstructed faces knowing the light directions and the one without for the first dataset.

In more details for the first dataset the results are summarized in Table 2 and examples of the reconstructed surfaces and the height maps are shown in figure 4. Also, the error difference between the reconstructed faces knowing the light directions and the one without is shown in figure 5. The same experiments were also performed for the second dataset and the obtained results are shown in Table 3 and examples of the obtained surfaces are shown in figures 6 and 7. In more details, for each dataset two tests were performed and in each test scenario three or four different images were used as input to the reconstruction system. The obtained height map in each scenario was compared with the equivalent height map obtained by the same input images but knowing the illumination directions.




Fig. 6. Reconstructed surfaces using the proposed method with unknown illumination on the right and the ground truth on the left under different view points for the second dataset.



Fig. 7. The error difference between the reconstructed faces knowing the light directions and the one without for the second dataset.




Fig. 8. Examples of profile views used to calculate the Hausdorff distance.

The proposed algorithm was further evaluated using the the Hausdorff distance comparing the surfaces obtained with and without any illumination information. In figure 8 results of the reconstructed faces obtained from the two cases are shown. Observing the results it can be inferred that the proposed methodology results very accurate estimates without knowing the light directions. In particular, the side view was used to evaluate the performance of the proposed approach. The background was extracted manually and the Hausdorff distance was used to compare the reconstructions with the original profiles. Table 4 shows the average results for all the faces.

5 Conclusions

In this paper, a method for 3D face reconstruction using photometric stereo with unknowing lights was proposed. A dimensionality reduction method was introduced to represent the face deformations due to the illumination variations and the self shadows in a lower space. The obtained mapping function was used to determine the illumination direction of each input image and that direction was used to apply photometric stereo. Experiments with faces were performed in order to evaluate the performance of the proposed scheme in a comparative study. From the experiments it was shown that the proposed approach results very accurate 3D surfaces without knowing the light directions and with a very small differences compared to the case of known directions. As a result the proposed approach is more general and requires less restrictions and information for the acquisition environment, allowing further applications on 3D face recognition and tracking.

Table 2. The accuracy of the proposed photometric stereo method for faces with unknown illumination directions in terms of height map percentage error over the ground truth (case of known directions) for the first dataset.

Average Height Error	Scenario A	Scenario B
Proposed Method	1.4974%	1.1161%
Schindler [25]	2.1775%	1.5553%

Table 3. The accuracy of the proposed photometric stereo method for faces with unknown illumination directions in terms of height map percentage error over the ground truth (case of known directions) for the second dataset [3].

Average Height Error	Scenario A	Scenario B
Proposed Method	3.0963%	2.7450%
Schindler [25]	3.7831%	2.9806%

References

1. N.G. Aldrin, S.P. Mallick, and D.J. Kriegman. Resolving the generalized bas-relief ambiguity by entropy minimization. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–7, June 2007.
2. Vasileios Argyriou, Stefanos Zafeiriou, Barbara Villarini, and Maria Petrou. A sparse representation method for determining the optimal illumination directions in photometric stereo. *Signal Processing*, 93(11):3027 – 3038, 2013.
3. G.A. Atkinson, M.F. Hansen, W.A.P. Smith, V. Argyriou, M. Petrou, M.L. Smith, and L.N. Smith. Face recognition and verification using photometric stereo: The photoface database and a comprehensive evaluation. *IEEE Transactions on Information Forensics and Security*, 2013.
4. Ronen Basri, David Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *IJCV*, 72(3):239–257, 2007.
5. P.N. Belhumeur, J.P. Hespanha, and D.J. Kriegman. Eigenfaces vs. fisherfaces: recognition using class-specific linear projection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19(7):711–720, 1997.
6. M. Belkin and P. Nivogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *NISP 14*, 2001.
7. M.K. Chandraker, S. Agarwal, and D.J. Kriegman. Shadowcuts: Photometric stereo with shadows. *CVPR*, June 2007.
8. A. Elgammal and C. Lee. Body pose tracking from uncalibrated camera using supervised manifold learning. *NIPS EHuM Workshop*, 2006.
9. A. Georghiadis. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. *9th ICCV*, 2, 2003.
10. P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, 1994.
11. X. He, S. Yan, Y. Hu, P. Niyogi, and H. Zhang. Face recognition using laplacian-faces. *IEEE Trans PAMI*, 27:328–340, 2005.
12. J. Hong and K. Song. Facial expression recognition under illumination variation. In *IEEE Workshop on Advanced Robotics and Its Social Impacts*, pages 1–6, 2007.
13. V. Silva J. Tenenbaum and J. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

Table 4. The average Hausdorff distance was used to compare the reconstructed with the original profiles.

Average Hausdorff Distance	Dataset 1	Dataset 2
Proposed Method	0.0097	0.0055
Schindler [25]	0.0124	0.0083

14. D. Fleet J. Wang and A. Hertzmann. Gaussian process dynamical models. *NISP 18*, 2006.
15. I.A. Kakadiaris, G. Passalis, G. Toderici, M.N. Murtuza, Y. Lu, N. Karampatziakis, and T. Theoharis. Three-dimensional face recognition in the presence of facial expressions: An annotated deformable model approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(4):640–649, April 2007.
16. S.N. Kautkar, G.A. Atkinson, and M.L. Smith. Face recognition in 2d and 2.5d using ridgelets and photometric stereo. *Pattern Recognition*, 45:3317–3327, 2012.
17. I. Kemelmacher-Shlizerman. Internet based morphable model. In *2013 IEEE International Conference on Computer Vision*, pages 3256–3263, Dec 2013.
18. I. Kemelmacher-Shlizerman and S. M. Seitz. Face reconstruction in the wild. In *2011 International Conference on Computer Vision*, pages 1746–1753, Nov 2011.
19. N. Lawrence. Gaussian process latent variable models for visualisation of high dimensional data. *NISP 16*, 2004.
20. KC. Lee. Acquiring linear subspaces for face recognition under variable lighting. *IEEE PAMI*, 2005.
21. F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 1490–1497, June 2013.
22. Thoma Papadhimetri and Paolo Favaro. A closed-form, consistent and robust solution to uncalibrated photometric stereo via local diffuse reflectance maxima. *International Journal of Computer Vision*, 107(2):139–154, 2014.
23. A. Robles-Kelly and E.R. Hancock. A graph-spectral approach to shapefrom-shading. *IEEE Trans. Image Process.*, 13(7), 2004.
24. S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
25. G. Schindler. Photometric stereo via computer screen lighting for real-time surface reconstruction. In *Proc. 3DPVT’08*, 2008.
26. B. Shi, Y. Matsushita, Y. Wei, C. Xu, and P. Tan. Self-calibrating photometric stereo. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1118–1125, June 2010.
27. B. Shi, Z. Wu, Z. Mo, D. Duan, S. Yeung, and P. Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. *IEEE Conference on CVPR*, pages 3707–3716, 2016.
28. G van Der Maaten, L ; Hinton. Visualizing data using t-sne. *Journal Of Machine Learning Research*, 9:2579–2605, November 2008.
29. V.Argyriou and M. Petrou. Recursive photometric stereo when multiple shadows and highlights are present. *Proceedings of CVPR*, 2008.
30. R. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
31. Z. Wu and P. Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 1498–1505, June 2013.